Beyond Free Riding: Quality of Indicators for Assessing Participation in Information Sharing for Threat Intelligence*

Omar Al-Ibrahim
University at Buffalo

Aziz Mohaisen
University at Buffalo

Charles Kamhoua
Air Force Research Lab

Kevin Kwiat
Air Force Research Lab

Laurent Njilla
Air Force Research Lab

arXiv:1702.00552v1 [cs.CR] 2 Feb 2017

ABSTRACT

Threat intelligence sharing has become a growing concept, whereby entities can exchange patterns of threats with each other, in the form of indicators, to a community of trust for threat analysis and incident response. However, sharing threat-related information have posed various risks to an organization that pertains to its security, privacy, and competitiveness. Given the coinciding benefits and risks of threat information sharing, some entities have adopted an elusive behavior of "free-riding" so that they can acquire the benefits of sharing without contributing much to the community. So far, understanding the effectiveness of sharing has been viewed from the perspective of the amount of information exchanged as opposed to its quality. In this paper, we introduce the notion of quality of indicators (QoI) for the assessment of the level of contribution by participants in information sharing for threat intelligence. We exemplify this notion through various metrics, including correctness, relevance, utility, and uniqueness of indicators. In order to realize the notion of QoI, we conducted an empirical study and taken a benchmark approach to define quality metrics, then we obtained a reference dataset and utilized tools from the machine learning literature for quality assessment. We compared these results against a model that only considers the volume of information as a metric for contribution, and unveiled various interesting observations, including the ability to spot low quality contributions that are synonym to free riding in threat information sharing.

1. INTRODUCTION

Today, the Internet connects millions of users, networks, and network collections worldwide, where the Internet's security and stability are quite important to the global economy and well-being of the human race. However, challenged by various forms of cyber-attacks, ensuring the security of the Internet and combatting the various attacks requires proper reconnaissance that prelude countermeasure actions. The information security and threat landscape has grown significantly, making it difficult for a single defender to defend against all of these attacks alone. As such information sharing for threat intelligence, a paradigm in which threat indicators are shared in a community of trust to facilitate defenses, has been on the rise [1].

In threat intelligence sharing, participants exchange patterns of threats with each others, in a form of threat indicators or signals. Participants are defined over a community of trust, and collaborate towards a common goal or mission; to understand and respond to emerging threats [2]. For such intelligence sharing to happen, standards for representation, exchange, and consumption of indicators are proposed in the literature [3, 4, 5]. Communities of trust are established, and systems and initiatives for sharing are built. For such initiatives to work, participants need to contribute information in those systems to be consumed by other community members. However, sharing threat-related information have posed various risks to organization, which pertain to security, privacy, and competitiveness. Given the coinciding benefits and risks of threat information sharing, some community members have adopted an elusive behavior of "free-riding" [6] so that they can achieve utility of the sharing paradigms without contributing much to the community.

So far, understanding the effectiveness of sharing has been viewed from the point of view of whether participants contribute or not using volume-based notions of contributions. Thus, a community member who does not contribute a volume of data (indicators) is a free-riding community member [7]. The state-of-the art on the problem did not include other metrics beyond simple measures of volume-based contribution, particularly metrics that capture and assess the Quality of Indicators (QoI) as a mean of understanding contribution in the information sharing paradigms.

1.1 Quality of Indicators

We believe that the nature of the information sharing as a concept and its application to threat intelligence both make quality a very relevant notion, and call for further investigation into the notion's definition and quantification in various settings. A well-defined measure of QoI could provide a better way of capturing contribution in general, and distinguishing community members who contribute useful data. Furthermore, threat intelligence systems present distinct challenges and opportunities to counter the problem of free-riding and other abusive behaviors once quality is defined. On the one hand, these systems generally lack the

*Approved for Public Release; Distribution Unlimited: 88ABW-2017-0416, Dated 01 Feb 2017. This work was done while O. Al-Ibrahim was visiting the University at Buffalo.

ACM ISBN 978-1-4503-2138-9.
DOI: 10.1145/1235

enforcement of a central authority, and therefore entities in these systems share information at their own will in a (somewhat) peer-to-peer fashion. As a result, it is necessary to envision a set of notions and mechanisms that characterize contribution in general, and are capable of capturing free-riding, while being implemented in a distributed manner and used by each community member. An ideal measure of QoI should be robust to distinguish between the various members based on their contribution, rather than a predefined notion of trust. With the possible speciality of community members, and the varying usefulness of indicators shared based on the context in which they are used, a major challenge is to assign context-dependent quality markers for indicators.

1.2 The Simple Contribution Measures

To the best of our knowledge, the problem of free-riding in information sharing for threat intelligence, while sparsely mentioned in other work [2, 8], is not treated properly in the literature. Thus, this work is the first of its type to be dedicated to the problem by identifying QoI as a new metric of contribution to capture free-riding in information sharing for threat intelligence.

Given the aforementioned challenges, one simple measure of contribution in information sharing systems weigh the volume of indicators contributed by various community members. However, for many reasons, such measure is insufficient, as described earlier. Therefore, it is important to understand the quality of shared information as a form of participation. Without a high quality of shared information, we cannot achieve actionable intelligence that is effective in combating cyber threats. Unfortunately, this issue is not well understood in the literature, and requires further exploration by identifying the meaning of quality, and basic methods and tools for assessing them are lacking.

1.3 Features of Quality of Indicators

In [1], Mohaisen *et al.* explored the potential correlation between QoI and privacy. However, privacy is not the only factor that affects QoI. One feature of quality is the *correctness* of an indicator; a meaningful annotation and label of the indicator that is true and accurate. A second possible feature of quality is the *relevance* of the indicator to the community members; because of the targeted nature of modern cyberthreats, information that is shared has to be contextual to the domain. A third plausible feature of quality of an indicator is its *utility*; informally, some indicators are more indicative than others about cyber-attacks, and therefore it is critical that participants in the threat intelligence community share information that capture prominent features of cyber-threats. Finally, the *uniqueness* of an indicator is another assessor of quality, which is defined as a measure of (dis)similarity with previously seen indicators. This property ensures that participants deliver indicators that are not duplicates or redundant, and provide additional threat information to other community members.

Besides these features, indicators are often time-sensitive, making temporal features very important when evaluating QoI. A timely indicator such as a source of an attack could be used to defend against an emerging attack, unlike a stale indicator that could be (potentially) used for postmortem analysis. As mentioned above, there is also potential correlation between QoI and privacy. Privacy can affect QoI (although perhaps negatively when privacy of indicator is ensured). We elaborate on this quality metric in this work, and show its quantification through data-driven analysis.

While each of the aforementioned measures can be used as a separate feature of quality, we envision that a single indicator could have multiple of those features. As such, we also assess QoI with respect to these metrics in the form of a weighted (continuous) score. Our method for evaluating QoI is based on exploiting a fine-grained historical records as benchmark for assessing the contributions of community members. We illustrate the concept through a concrete evaluation of a real dataset from various security vendors of antivirus scans and their results of labeling malware samples as seen in the VirusTotal service (https://virustotal.com/).

1.4 Contributions

The main contributions of this paper are multifold. First, we identify the need for QoI to capture contributions by community members in information sharing paradigm. QoI captures a wide spectrum of behaviors, from altruistic behavior, where a community member contributes a lot of (high quality) indicators to free-riding, where a community member contributes less, or contribute a lot of low quality indicators. Second, we develop and formulate various metrics that are robust to capture the notion of quality. Third, we experimentally demonstrate those measures and metrics, and show their robustness, and how they differ in identifying contributor's behavior (particularly free-riding) from the simple volume-based measure of contribution.

1.5 Organization

The organization of this paper is as follows. In Section 2, we provide an overview on cyber-threat intelligence and the risks of information sharing. In Section 3, we provide an overview of our quality of indicator(QoI) assessment methodology. In Section 4, we describe the processes involved in our QoI-based assessment. In Section 5, we present the results of our benchmark experiment, afterwards we discuss related work in Section 6 and finally we conclude in Section 7.

2. OVERVIEW AND PRELIMINARIES

We first provide an overview of cyber-threat intelligence systems, then introduce to unique problem with information sharing in these systems which demand quality measures.

2.1 The Threat Landscape

The Internet today connects hundreds of millions of users worldwide, and is operated by service providers who connect businesses, education institutes, and government agencies, collectively forming a global village. In the recent years, the Internet has been challenged by various forms of cyber attacks, ranging from endpoint malware attacks [9] to massive network disruptions and instabilities [10].

At the endpoint side, malware is capable of penetrating a perimeter's security in many enterprise systems, exfiltrating sensitive data from such systems, and causing great damage to both private and public sector networks [11]. At the larger scale, multiple endpoint infections by malware are more powerful, and pose a greater risk, seen often in systemized large-scale botnets [12]. Botnets, defined as collections of networks of infected hosts are the basic fabric for the operation of many cybercriminal activities. Botnets rely on principled designs, where bots (infected hosts in a botnet)

execute commands on behalf of their herder (botmaster), utilizing command and control (C&C) infrastructure [13].

Botnets today are used for a variety of cybercriminal activities, including spam, massive denial-of-service (DDoS) attacks, and data exfiltration, among many others. Botnets represent a major component of the cybercrime ecosystem, with the rise of botnet-as-a-service. Today, hackers utilize network reconnaissance to probe targets for vulnerabilities and craft custom payloads to gain control over their infrastructure by spreading malware in propagation efforts.

2.2 The Need for Threat Intelligence

Defending against the threat vectors of malware and botnets is a challenging task, which resulted in a rich body of literature. The literature on defending against malware and botnets looks into identifying "signals", "indicators", or simply "features" that could be useful in identifying endpoint systems, malware, and botnets. For malware, for example, such features could include static strings in that piece of malware, dynamic artifacts that the malware generates when executed in the wild, or external context information associated with the binary binary of the malware (such as the author's information, operating system, etc.). For botnets, the C&C infrastructure may include domain names and Internet Protocol (IP) addresses, and knowing such information can be very helpful in identifying a botnet. For example, botnets tend to use Domain Generation Algorithms (DGAs) [14], which result in random domain strings with high entropy, and being able to identify those domain names is key to detection of botnets. Furthermore, being able to distinguish between various DGAs is key to attribution of threat to a certain botnet family. An effective cyber defense would rely on a good visibility into many of those features.

Combatting cyber threats and attacks requires intelligence gathering that prelude countermeasure actions, as seen in the above examples. To this end, cyberthreat intelligence has become a growing concept. Today, organizations in the public and private sector, government and industry, have established tools seeking first-hand knowledge about new cyber-attacks and malware threats. This includes the ability to recognize and act upon indicators of attack and compromise scenarios, essentially putting the pieces together for analysis about attack methods and processes using static and dynamic analysis and profiling techniques, open source, social media, and dark web intelligence.

2.3 Threat Intelligence Sharing

The need for information sharing for threat intelligence is necessitated by both economical and technical realities. Being able to identify all the types of indicators and features useful for characterizing, identifying, and defending against all types of threats, while desirable, is infeasible from both technical and economical standpoints. With new technologies such as cloud, mobile computing, social networks, and the Internet of Things (IoTs), and the the persistence of adversaries through cybercrime and advanced persistent threats (APTs) have also brought several challenges. Therefore, it is reasonable to say that no single player in this ecosystem is capable of addressing all security issues alone.

For this reason, sharing information of threat intelligence among vendors and government entities has emerged as a plausible technique for efficiently and effectively defending against new and emerging threats. With threat intelligence sharing, operational experience is communicated to other parties in an ecosystem to enable them to effectively defend against current attacks, and to improve their defense posture by preventing such attacks from happening utilizing such actionable intelligence.

To enable information sharing, organizations need to agree on standardizing threat information. This requires defining the content fields, encoding, and exchange format of the information relevant to a particular threat or incident, along with a pre-defined protocol to communicate the criticality of such information. Various standards for information sharing have been proposed [3, 4, 5] to automate and structure the exchange of threat information with a community of trust.

Today, standards are used in the exchange of indicators of software, hardware, and network artifacts, and are intended for operationalizing those indicators in many applications, including security operations related to malware characterization, vulnerability analysis, remediation, platform hardening, and incident-response [1].

2.4 Risks of Sharing

While threat information sharing brings many benefits to the sharing community members, it may incur security risks about participants, their operational contexts, and security posture. Not only that, the same information, once exposed to an adversary, may be used to test their applicability on other target systems, who may lag behind in security updates or miss out on patching vulnerabilities. Therefore, the adversary will be able to utilize such information for attacking other unpatched systems.

The risk of sharing may go beyond fingerprinting systems to leaking personal identifiable information about individuals. Various types of sharing standards are proven to leak personal identifiable information (PII) that may contain names, email addresses, and other types of sensitive data [15]. For example, privacy violations in sharing standards may occur in the form of a document which contains contact information for the constituent responsible for an incident report. This type of information may become personally identifiable in the case when the contact information of a particular individual are used.

Participants in a threat intelligence sharing community may interact with one another with various degrees of collaboration and competition, which may affect the way they share [8, 2]. Because of that, many companies and organizations today are reluctant in sharing firsthand intelligence, and mostly gather and ingest information from neighboring sources that are less significant.

2.5 Formulation of the Free-Riding Problem

Given the triad of security, privacy, and risks associated with threat intelligence sharing, some members might be joining communities of sharing for the purpose of benefiting from the platform without offering valuable information themselves, hence, the term, "free-riding" is coined to refer to the behavior of such users who act to maximize their own utility at the expense of the welfare of the community.

This problem is not new, and is manifested in other distributed settings, most notably Peer-to-Peer (P2P) systems. In P2P systems, cooperation is required for the operation the system. However, cooperation may incur significant communication and computational overhead, thus users may refuse to contribute their fair share of resources. At the same time

they may utilize the system by consuming the resources of other peers. Researchers demonstrated the impact of free-riders in P2P systems, such as BitTorent [16], and observed a significant increase in download times for high contributing nodes in the presence of few low contributing ones.

2.6 How Quality of Indicators Help

So far, and to the best of our knowledge, understanding the effectiveness of sharing has been viewed from the point of view of whether participants contribute or not (thus the literal meaning of free-riding). This form of contribution is perceived as a volume-based contribution, since the level of contribution by any participant is evaluated directly by the amount of information communicated to the community regardless of its nature, whether it is used by community members or not. Given the large amount of unprocessed threat-related events, which are generated by automation tools, such as security information and event management (SIEM) technologies, and the fact that in many sharing systems today, the contribution level amounts to the volume of data, actors may find it more convenient to submit raw, unprocessed, or unused events as indicators of threat to avoid the investment on resources for cleaning, contextualizing, operationalizing, and filtering such information. For this reason, it is important to consider the quality of shared information as the basis for evaluating the level of participation, because a simple and coarse measure of participation is insufficient. In order to overcome these obstacles, in this paper we propose QoI as a quantifiable and measurable metrics and provide a methods for quality assessment.

3. QOI ASSESSMENT METHODOLOGY

Assessing QoI is a nontrivial task. However, as other quality notions (e.g., quality of services, quality of experience, etc.), QoI requires defining metrics, methodology for assessing such metrics, and methods for validation of the proposed metrics based on sound assumptions, to find out what capabilities they provide, including addressing free-riding.

Our approach in developing QoI metrics is intuitive, and uses several sounds assumptions driven from the context in which indicators are used. In particular, QoI metrics include correctness of an indicator with respect to a label feature, the relevance of the indicator to a consuming community member, the utility of the indicator, and its uniqueness. More (informal) details are provided in section 3.1.

Our approach for assessment of QoI uses a reference "golden" dataset as ground truth. Based on whether the indicator provided by a community member is in the reference dataset or not, the assessor proceeds by either matching the metric attribute of the indicator to that of the golden dataset. If the indicator is not in the golden set, and assuming that the golden dataset has "similar indicators", a machine learning algorithm is used for predicting the attribute of the QoI metric, and compare it to the one provided by community member. For an arbitrary number of indicators provided by the community member, a score is then established for that community member based on the normalized weighted sum of QoI values across all indicators.

In this section, we describe the quality metrics and the details of our methodology including a system architecture for assessment, processes followed and the data-flow.

3.1 QoI Metrics

In the following, we identified four metrics to be used for assessment of quality: correctness, relevance, utility, and uniqueness, as described below.

3.1.1 Correctness

For a given reference dataset, the correctness metric of QoI captures whether attributes of an indicator (e.g., label used for attribution, severity score used for risk assessment, etc.) are consistent with the assessor's reference. For that, and using the labeling of a malware sample as an example of an attribute for an indicator, we compute the correctness score as the fraction of samples that match the anticipated labels. Specifically, this is computed as the aggregate binary score of the correct samples normalized by the size of the sample set of indicators.

3.1.2 Relevance

Informally, the relevance metric of a QoI measures the extent to which an indicator submitted by a community member to the community is contextual and of interest to the rest of the community. As such, in defining and assessing the relevance of an indicator, we use a reference weight assignment to the class labels giving higher weight to labels of greater interest to a particularly community member (assessor) and lower weights to less desirable labels. The relevance score is then computed as the average weighted sum for all sample indicators in the set.

3.1.3 Utility

The utility is similar to the relevance of an indicator, although at a finer-granularity than an indicator. As such, we view the utility of an indicator as the average weighted sum of all of its feature components. This is, we assign a different weight to each feature of the indicator to leverage features that are a better candidate input to prediction of threats. While the weighting of the features of an indicator could be realized using one of many ways, we suggest the information gain as a measure of weighting features. For example, using a similar notion, the weights of the feature components can be computed using various statistical models for sensitivity analysis, including the Principal Component Analysis (PCA) technique [17]).

3.1.4 Uniqueness

The uniqueness of an indicator is a measure of the (dis)-similarity of the indicator in comparison with other submitted indicators by contributors in the community. A vector distance (e.g., using the Mahalanobis distance [18], which captures the difference between a point $\vec{x}$ and a distribution of points $X = \{\vec{x_1}, \vec{x_2}, \ldots, \vec{x_n}\}$ with a mean $\vec{\mu}$) is computed to determine the degree of uniqueness. We also define a threshold on the minimum distance between the feature vector of an indicator to other indicators, and use that threshold to tell whether an indicator is unique or not.

3.2 System Architecture & Design

Having elaborated on the informal definition of QoI metrics, we now move to discuss the QoI system architecture, first as a strawman highlighting the main concept of assessing QoI, and then as a fully functioning system that addresses various issues in the strawman design.

Our system for assessing QoI operates for a set of distributed nodes in a community of trust. Those nodes are

are logically connected with each other in a P2P fashion, as shown in Figure 1. As such, each of these nodes would participate in the sharing and consumption of threat indicators provided by other peers, which is achieved within a community of trust that is separated from other communities in the sharing ecosystem. Before nodes can accept and operationalize (process) these threat indicators, they need to evaluate their quality by asking a special node, an assessor, which has sufficient information to perform such function, for a rating (scoring) of the indicator. In response, the assessor assigns a quality score for the indicator based on a ground-truth available to the assessor, and using a reference dataset the assessor has access to.

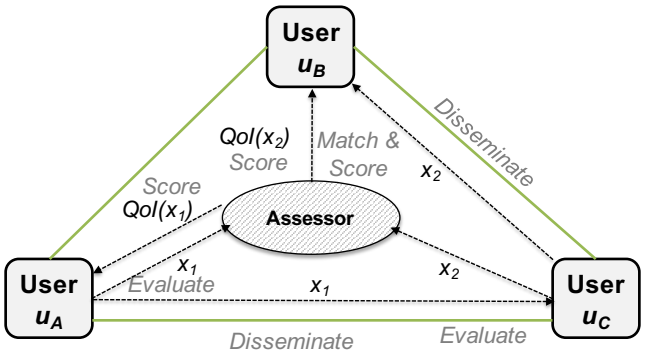

Figure 1: A strawman design for an architecture to assess QoI using a centralized assessor.

In this architecture, we assume that the messages sent between the peers in the system are authentic and tamper-evident, using existing threat exchange protocols that provide end-to-end security guarantees. One issue that the above architecture may suffer from is the amount of trust each community member has to put in the assessor, and the validity of his scoring of indicators. We address this issue by assuming that the architecture can support more than one assessor node, and these nodes may coordinate among each each other. We leave exploring the spectrum of the number of assessors and the rationale for various numbers as a feature work. However, as a feature, a consumer of a score by the assessors (where the assessors' number is greater than 2) could perform a majority voting to improve the robustness of the scoring to address intentional bias (dishonesty of assessors, if any), and unintentional bias (due to issues with the underlying ground truth dataset). On the other hand, any and every node in the system could act as an assessor, if the golden reference dataset is available to them.

While the latter assumption of the availability of the data to each node in the system is very implausible, a milder assumption for the operation of the strawman design above is the coverage of data: the system assumes the reference dataset has sufficient information about every possible indicators presented by the various community members. However, based on a prior work assessing coverage of indicators [19], no single community member (antivirus scanner) in the case of malware detection and labeling has a 100% coverage or accuracy. Based on the same study, and for a malware family such as Zeus [20], it takes 6 and 18 community members to provide close to perfect coverage of detection and correctness of labeling, respectively. Such numbers are close to 10% and 30%, respectively, of the entire set of community members with antivirus scans in the VirusTotal dataset. This in particular calls for a more "intelligent" process for the assessment of QoI, using not only explicitly provided labels, but also using learned labels from features of indicators utilizing advanced machine learning techniques.

3.3 System Setup and Steps

At a high-level, our ideal system for assessment has the following specific procedures for system setup:

1. *Defining quality metrics and scoring procedures.* Quality metrics are used as a measurement criteria to ensure that community members who participate in information sharing provide threat indicators that are valuable to other members, while scoring procedures are methods that specify how these metrics are used to generate a quality score.

2. *Defining annotations for threat and quality labeling.* Annotations can either be labels that indicate the type of threat or they can be labels for identifying the quality (severity, timeliness, etc.) level or quality type of an indicator. Utilizing these annotations, a weight value is assigned to each quality label, and a scoring method is utilized to convert the quality labels to a numeric aggregate score for the indicator.

3. *Building the reference dataset.* The reference dataset will be used to evaluate QoI for a sample of indicators submitted by a sample provider. To build the initial reference dataset, data that is collected through security operations (e.g., monitoring, profiling, analyses, etc.) is vetted for their validity and applicability to the domain, perhaps using often expensive by necessary manual vetting [19, 21].

4. *Defining extrapolation procedures and training the classifier.* Extrapolation procedures enable a quality assessor to predict the label of an indicator using its feature set and classifier model. The classifier is trained using a supervised learning process extracted from the reference dataset. This reference dataset is collected for the purpose of initializing the system.

After the initial setup of the system, the sample indicator is assessed for its quality and a quality score is computed. The following describes the steps of the assessment. 1) Obtain a set of sample indicators where each sample is composed of a tuple (label, vector) that consists of a label and a vector of features. 2) For each sample, extract the feature vector and feed the data as test input to the trained classifier which predicts its label. 3) Compare between the predicted label and the label provided by the sample. Indicate whether the two labels match, and record the comparison result as a quality annotation. 4) Compute the confidence level and include other quality annotations for the indicator using the labeling rubric. 5) Use scoring procedures and quality labels to compute a quality score for the indicator.

3.4 QoI Assessment Process Operation

The illustration for the complete process and dataflow of the QoI assessment method embraced in our design is depicted in Figure 2. As can be seen in the figure, the setup of the system is achieved through the use of a supervised

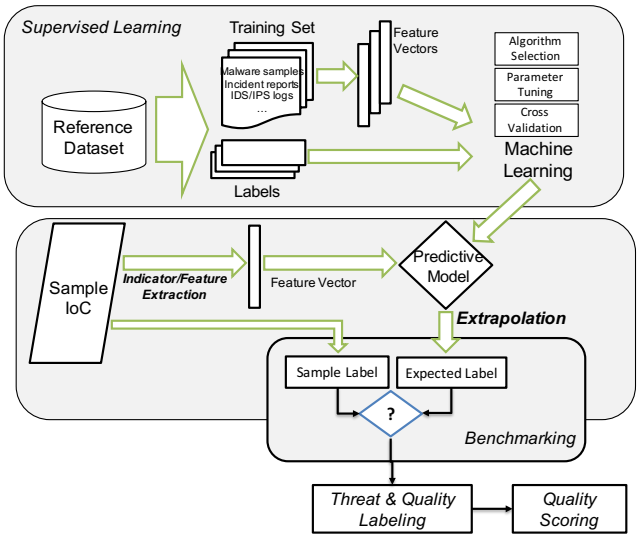

Figure 2: The QoI assessment process, incorporating a reference model established using the training procedure of a machine learning algorithm and predicted label of an indicator provided by a community member. For correctness of indicators as the target metric, a high quality is achieved when the predicted label matches the label provided by the community member (sample IoC label).

learning process over the reference dataset, rather than the direct matching of explicit labels of indicators in the dataset.

In this assessment system, we assume a reference labeled (training) dataset that contains a comprehensive library of artifacts, such as malware samples, incident reports, and logs, and that has been collected through typical operational intelligence gathering procedures. Ways for obtaining such labels falls out of the scope of this work, and industry's best practices, as described in [19] could be used. Upon ingesting those artifacts in our system, they are converted into a set of training samples with their corresponding features. Each sample is a pair that consists of a feature vector as an input object for the machine learning algorithm and the corresponding threat label as the desired output value for the object class. To reach such end goal of predicting a label correctly, the build of our trained model encompass multiple components, namely a feature selection procedure, a machine learning algorithm selection procedure (e.g., SVM, logistic regression, random forest, etc.) and the corresponding parameters (e.g., procedure for regularization and linearization in case of SVM and LR, respectively), and cross validation procedures (e.g., fold size, validation strategy, etc.).

Upon building the machine learning model, establishing a confidence in its performance through typical evaluation metrics (e.g., low false positive and false negative, and high true positive and true negative), we then use the built model as a predictor for previously unseen threat indicators. In such operational setting, given a sample indicator provided by a community member, and before ingesting it the community member would pass it to the assessor for further evaluation and scoring. The assessor then extracts a feature set corresponding to the indicator using a standard form, and converts it into a feature vector. The assessor then uses the previously built model as a predictor, and assigns a label (e.g., using a multiclass SVM, the assessor can indicate the label closest in the training set to that of the newly observed indicator). The assessor then decides the quality of the indicator by taking both the predicted and the self-provided label by the community member into account. The quality scoring engine then uses the individual scores of the various indicators provided by each community member to assess their actual contribution, and detect free-riders.

We note that the "intelligent" system above addresses various issues in the strawman system. First, rather than requiring the actual indicators to be present in the ground-truth dataset, this technique requires only the availability of a sufficient number of indicators of the same label. Second, with such flexibility in defining the ground-truth through a learning and model-building process, the number of community members that can act as assessors greatly increase. Finally, the even when the labels of indicators are not provided by the community member, e.g., artifacts provided to the community are not operationalized, this technique makes use of those indicators through other measures of QoI, such as relevance, or utility, or uniqueness, which do not require a label to be provided by the contributing community member.

4. QOI ASSESSMENT PROCEDURES

As discussed in section 3.3, the QoI assessment process is composed of a series of steps in order to initialize and operate our system for both assessment of individual indicators and scoring of community members as a whole. Specifically, these steps begin with collecting the reference dataset and building the prediction model, then extrapolating, benchmarking, and computing a quality score for a given indicator. In the following, we outline in more details each of those processes and procedures.

4.1 Reference Dataset and Learning

After identifying metrics for defining quality, as exemplified in section 3.1, we demonstrated the use of QoI for the assessment of the contribution level of participants. As mentioned before, our methods for computing the QoI involve multiple processes. In order to initialize the system the reference dataset is used to build prediction model through supervised techniques of machine learning.

Specifically, this involves submission of sample artifacts from multiple sources, and the premise is to utilize the notion of quality as opposed to a simple view of contribution based on the volume of data (i.e. the number of samples). In order to evaluate the QoI provided by community members, a reference dataset is used as a resource of ground truth.

While proposing methods for obtaining ground truth falls out of the scope of this work, we only use an example in this work to bootstrap the evaluation. In particular, in this work we demonstrate evaluating the quality of malware labels by AV vendor using VirusTotal as a reference dataset with samples that are manually vetted by one community member [19] (in such setting VirusTotal could be loosely defined as the community of trust). In short, VirusTotal is a multi-engine AV scanner that accepts submissions by users and scans the samples with those engines. The results from VirusTotal provide many useful artifacts and annotations, including the labeling of a sample by the various engines of AV scanners, as well as other behavioral and static features of malware samples. Though there might be some inconsis-

tencies in the final labeling between the results of VirusTotal and across vendors, the premise is that the tool can be trusted for samples that have been submitted multiple times over sufficient periods of time, particularly since AV vendors update their results with VirusTotal whenever acquiring a new signature for a previously unknown sample.

Features provided in the ground truth are particularly useful in a learning a model for their label prediction. To identify the family to which a malware sample belongs, security vendors (AV scanners, community members) usually gather various characteristics and features of the sample using static analysis, dynamic analysis, and memory forensics. For static analysis, artifacts like file name, size, hashes, magic literals, compression artifacts, date, source, author, file type and portal executable (PE) header, sections, imports, import hash, among others, are used. For dynamic analysis, file system, user memory, registry, and network artifacts and features are collected. For memory forensics, memory byte patterns are captured to create a signature.

4.2 Extrapolation and Benchmarking

After building the prediction function by training classifier with the reference dataset, the next step is to assess QoI through extrapolation from the prediction function results. The remaining question becomes how the reference set are used to assess and extrapolate the values and quality of indicators. In order to answer this question, we elaborate on a particular machine learning techniques, the semi-supervised learning and its application to the problem at hand.

Classifier model. While our system described in the previous section uses multiple off-the-shelf algorithms, we highlight the operation of QoI using a classifier model based on the nearest centroid classifier (ncc), specifically we adopt a variant called the linear discriminant analysis [22] to map threat indicators to their respective labels. In this model, each label is characterized by its vector of average feature values (i.e. class centroid). A new sample indicator is evaluated by computing the scaled distance between the features of the sample and each class centroid, and then the sample is assigned to the class to which it is nearest.

To build the classifier, we obtain r samples for training from the reference dataset. This dataset is built such that there are r_i training samples per class, with d features per sample. For each training sample y, we observe a label $\ell \in \Lambda$ and a sample vector $\vec{y}$. For simplicity we refer to the classes labels by their indices $i = 1, 2, \ldots, \lambda$. Note that each $\vec{y}$ is a vector of length m. We assume that samples labeled by i are distributed as $\mathcal{N}(\mu_i, \Sigma)$, the multivariate normal distribution with mean vector μ_i and standard deviation matrix Σ. We denote by $\mathcal{L}(x, \mu_i, \Sigma)$, the corresponding probability density function. Finally, let π_i be the prior probability that an unknown sample comes from class labeled by i.

Bayes' Theorem states that the probability that an observed sample x comes from class i is proportional to the product of the class density and prior probability:

$$P(Z = i|X = x) \propto \mathcal{L}(x, \mu_i, \Sigma) \times \pi_i \quad (1)$$

where $P(Z = i|X = x)$ is the posterior probability that sample x comes from class i. The classifier assigns the sample to the class with the largest posterior probability to minimize the misclassification error. This can be written as a rule:

$$\hat{z}(x) = \arg\min_i \{(x - \mu_i)^T \Sigma^{-1} (x - \mu_i) - 2\log(\pi_i)\}. \quad (2)$$

Therefore, a sample is assigned to the nearest class and the distance is computed using the LDA metric: $||x - \mu_i||^2_\Sigma - 2\log(\pi_i)$, where $||x-\mu||^2 = (x-\mu)^T\Sigma^{-1}(x-\mu)$ is the square of the Mahalanobis distance between x and μ.

Misclassification rate. A misclassfication occurs when an indicator is assigned to an incorrect label. The probability of making a classification error $P(\epsilon)$ is:

$$P(\epsilon) = \sum_{j=1}^{\lambda} [P(\hat{Z} \neq j|Z = j) \times \pi_j]. \quad (3)$$

The misclassification rate using the LDA rule can be derived from (2). In particular, we can calculate the misclassification rate of the nearest-centroid using:

$$P(\epsilon) = \sum_{j=1}^{\lambda} [1 - \phi(\min_{i \neq j} \left\{ \frac{\|\mu_j - \mu_i\|^2_\Sigma + 2\log(\frac{\pi_j}{\pi_i})}{2\|\mu_j - \mu_i\|_\Sigma} \right\})] \times \pi_j, \quad (4)$$

where ϕ is the cumulative distribution function (cdf) of the standard normal distribution. Note that this assumes that the sample data are normally distributed as stated by the model. The equation above can be interpreted as a measure of the collective distance between all of the class centroids taking into account class prior probabilities. In general, the misclassification rate is small when the centroids are far apart and increases otherwise.

4.3 Labeling and Quality Scoring

We used the ncc to predict labels for observed indicators and compared the results with the sample labels. This enables us to compute a score on the correctness and quality of the feature set for the indicators. In the following we formalize the steps to compute a score for the samples based on the quality metrics described earlier.

Denote by n the number of users in the system. Each user u_i provides a set of samples $X_i = \{(\vec{x}_{i1}, l_{i1}), (\vec{x}_{i2}, l_{i2}), (\vec{x}_{i3}, l_{i3}), \ldots, (\vec{x}_{ik}, l_{ik})\}$ with feature vector $\vec{x}_{ij}$ and sample label $l_{ij} \in \Lambda$ for $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, k$.

4.3.1 Correctness

As described earlier, the reference dataset is used as the benchmark for determining the correct label for an arbitrary sample. Each sample consists of a feature vector and an associated label. In procedure 1 we outline the algorithm used for computing the correctness as a QoI metric.

As shown above, we first build a classifier by utilizing the reference dataset Y as the training set and forming a prediction on the label of $\vec{x}$ to obtain l'. Then, the assigned label of $\vec{x}$ is compared against the predicted label l' and a positive score is given if labels match. The correctness is computed as the average sum of scores for all samples in X_i.

4.3.2 Relevance

The steps for computing the relevance variable of a set of indicators are shown in Scoring Procedure 2. As can be seen, the weight values $\omega_{r_1}, \omega_{r_2}, \ldots \omega_{r_{|\Lambda|}}$ are arbitrarily chosen, and a mapping function $w_R(.)$ is defined to assign weights labels such that higher weight values are assigned to labels of greater interest to the community members.

For each sample x, the corresponding label is evaluated using the mapping function $w_R(.)$ to obtain the weight value as the sample score. The relevance score of X_i, denoted by $R(X_i)$, is calculated as the average weight sum of the scores.

Scoring Procedure 1 Correctness of X_i (C)

1: Obtain reference dataset $Y = \{(\vec{y}_1, l_1), (\vec{y}_2, l_2), (\vec{y}_3, l_3), \ldots, (\vec{y}_r, l_r)\}$ where $l_i \in \Lambda$ for $i = 1, 2, \ldots, r$.
2: Evaluate X_i by applying the ncc method as follows:

a *Training*: Compute reference indicators $\vec{\mu}_1, \vec{\mu}_2, \ldots, \vec{\mu}_{|\Lambda|}$ for class labels in Λ, as per-class centroids $\vec{\mu}_\ell = 1/|Y_t| \sum_{(\vec{y}_i, l_i) \in Y_\ell} \vec{y}_i$, where Y_ℓ is a subset of Y belonging to the class label $\ell \in \Lambda$.

b *Prediction*: For every sample $\vec{x}_{ij}$, apply the classifier function to compute the label,

$$l' = \arg\min_{l \in \Lambda} ||\vec{\mu}_l - \vec{x}||$$

3: For every sample, compute the sample score $s_c(\vec{x}_{ij})$ as:

$$s_c(\vec{x}_{ij}) = \begin{cases} 1 & l'_j = l_{ij} \\ 0 & \text{otherwise} \end{cases}$$

4: Compute the correctness score (C) of X_i by taking the average sum: $C(X_i) = \frac{1}{k} \sum_{j=1}^{k} s(\vec{x}_{ij})$

Scoring Procedure 2 Relevance of X_i (R)

1: Define weight values: $\omega_{r_1}, \omega_{r_2}, \omega_{r_3}, \ldots, \omega_{r_{|\Lambda|}} \in \mathbb{R}$
2: Define a weight function $w_R(.)$ to assign elements in the label set $l_i \in \Lambda$ as $w_R(l_i) = \omega_{r_i}$
3: Compute the relevance of X_i as average weighted sum:

$$R(X_i) = \Big(\sum_{(\vec{x}_{ij}, l_{ij}) \in X_i} w_R(l_{ij}) \Big) / \Big(\sum_{k=1}^{|\Lambda|} \omega_{r_k} \Big)$$

Scoring Procedure 3 Utility of X_i (U)

1: Define utility types $t_1, t_2, \ldots, t_d \in \mathbb{R}$
2: Define weight values $\omega_{t_1}, \omega_{t_2} \ldots \omega_{t_d}$, where each weight value corresponds to a utility type.
3: Define a weight function $w_T(.)$ s.t. $\ell \in \Lambda$ maps to a utility weight, i.e. $w_T(\ell) = \omega_{t_m}$ where $m = \{1, 2, \ldots, d\}$.
4: **for** $x \in X_i$ **do**
5: Compute a weight of $x = (\vec{x}, l')$, using $w_T(l') = \omega_{t_{l'}}$
6: **end for**
7: Compute the utility score of X_i as the average sum of the sample weights: $U(X_i) = \frac{1}{k} \sum_{j=1}^{k} \omega_{t_j}$, where t_j is the corresponding label type of sample $x_{ij} \in X_i$

Scoring Procedure 4 Uniqueness of X_i (N)

1: Consider the set Z which is initially empty, i.e. $Z = \phi$
2: Build the set Z by considering unique samples from the sets $X_1, X_2, ...X_n$
3: **for** $i = 1, 2, \ldots n$ **do**
4: **for** $j = 1, 2, \ldots k$ **do**
5: **if** $x_{ij} \notin Z$ **then** add x_{ij} to Z.
6: **end if**
7: **end for**
8: **end for**
9: Compute the uniqueness score for X_i as follows

$$s_n(x_{ij}) = \begin{cases} 1 & x_{ij} \in Z \backslash \{X_i\} \\ 0 & otherwise \end{cases} \quad (5)$$

10: Compute the uniqueness score (N) of X_i by accumulating the scores of the elements: $N(X_i) = \frac{1}{k} \sum_{j=1}^{k} s_n(x_{ij})$

4.3.3 Utility

Next, we provide the sequence of steps required for calculating the utility variable of a set of indicators, in procedure 3. In this procedure we note that the utility of an indicator is determined by the sum of the utility weights of the samples. The weights $\omega_{t_1}, \omega_{t_2} \ldots \omega_{t_d}$ and weight function $w_T(.)$ are defined by the application.

4.3.4 Uniqueness

Another metric of QoI is their uniqueness, where highly unique indicators are considered more valuable than common indicators. In procedure 4 we outline the steps used for calculating the uniqueness of a set of indicators.

In this procedure, we assume that samples can be uniquely identified (e.g. using hashes). In set notation, we can say that an element $x_{ij} \in X_i$ is unique if it is not an element of other sample sets, i.e. $x_{ij} \notin \bigcup X \backslash \{X_i\}$.

4.3.5 Quality of Indicator (QoI)

QoI is a comprehensive measure of the various notions of quality defined earlier. In particular, QoI for X_i is the average weighted sum of the four components: correctness (C), relevance (R), utility (U) and uniqueness (N), as shown in procedure 5. The weights assigned for individual metrics are application- and community member-specific.

5. EVALUATION AND FINDINGS

In this section, we evaluate the scoring method for contribution based on quality of indicators, and highlight how it addresses the free-riding problem in information sharing in a unique way. We start by analyzing the dataset that we obtained from AV vendors about their sample labeling, then we utilize this dataset and apply QoI-based and volume-based scoring methods to compare between the vendors.

5.1 Dataset Characteristics

To highlight QoI as a new notion of evaluating contribution in information sharing for threat intelligence, we compare the difference between quality-based and volume-based scoring methods for the contribution of AV vendors. To this end, our dataset enumerates AV vendors who submitted their artifacts of malware samples, including labels, to VirusTotal during the period of our data collection from mid 2011 to mid 2013 [19]. A key goal of the evaluation is to demonstrate the deficiency in the use of volume-based scores, since one vendor can achieve a high rating by submitting a large number of artifacts about malware samples of low quality. As discussed previously, this could happen because of several reasons: the submitted artifacts about some malware sample are incorrect, the sample family is uninteresting, or that the kind of information submitted about the samples are not helpful in identifying or detecting them.

Table 1 depicts the malware families used in this study, their sample size, and the corresponding brief description of each family type. All scans are carried out on those malware samples around May 2013 timeframe. The dataset provides a diverse representation of families, which nicely facilitate

Scoring Procedure 5 QoI of X_i (QoI)

1: Define normalized weights for the components: $\omega_C, \omega_R, \omega_U$, and ω_N.
2: Calculate the quality of indicator (QoI) as the weighted sum of the components: $QoI(X_i) = \omega_C C(X_i) + \omega_R R(X_i) + \omega_U U(X_i) + \omega_N N(X_i)$

Table 1: Malware families used in the study. DDoS stands for distributed denial of service. Also, Ddoser is known as BlackEnergy while Darkness is known as Optima. Dataset and description are from [19].

Malware family	#	Description
Avzhan	3458	Commercial DDoS bot
Darkness	1878	Commercial DDoS bot
Ddoser	502	Commercial DDoS bot
jkddos	333	Comercial DDoS bot
N0ise	431	Commerical DDoS bot
ShadyRAT	1287	targeted gov and corps
DNSCalc	403	targeted US defense companies
Lurid	399	initially targeted NGOs
Getkys	953	targets medical sector
ZeroAccess	568	Rootkit, monetized by click-fraud
Zeus	1975	Banking, targets credentials

our study. As can be seen, Avzhan and Darkness are the most popular DDoS malware being submitted. On the other hand, ShadyRAT is the most popular targeted malware with the largest sample size in its category (represents about 43% of the targeted samples and 24% of the total samples), while Zeus has the largest sample size for Trojan malware, roughly about 77% of the Trojan samples and 16% of the total samples. Furthermore, we observe that more than half of the samples are DDoS (54%), 21% are Trojans, and 25% are targeted malware.This breakdown provides an insight about the threat landscape and the frequency in which these types of malware appear in the wild. Oftentimes AV vendors harvest malware samples by utilizing deployment of Internet sensors for packet capturing, or using isolated environments such as honeypots and virtualization tools for behavioral analysis. The increased number of samples for DDoS is justified by the need for vast deployment for scaling up the number of infected hosts to launching attacks. On the other hand, targeted malware are less common in the wild because they are deployed in limited number of hosts, and are typically designed with covertness in mind. In the rest of this analysis, the identity of the vendors is anonymized.

5.2 Results

In this section, we introduce the results and finding by performing an evaluation of the various QoI metrics over our evaluation dataset (we evaluate all but uniqueness, since it is trivial to assess). First, we note that while there are more samples gathered for DDoS-type malware in comparison with others, the threat-intelligence community often gives more weight to identify malware or incidents that are less observable, which present a level of sophistication. Thus, for our evaluation, we consider trojan and targeted malware more relevant than DDoS, from the point of view of community members consuming the shared information.

Relevance. Figure 3 presents the normalized scores for the relevance of indicators for each vendor. In assessing relevance, we give more weight to targeted malware and Trojans over DDoS samples of each vendor (community member). Specifically, the weights are $\omega_{targeted} = 5$, $\omega_{troj} = 3$, and $\omega_{ddos} = 1$m and "0" otherwise. As shown, in the relevance-based assessment a higher score is given to vendors who contribute more targeted and Trojan samples, and vendors who only contribute DDoS samples are greatly de-emphasized. We observe two distinct behaviors. First certain contributors who have high volume-based score tend to have a very low (close to "0") score when using the QoI metric of relevance for their evaluation. In particular, with the two relevant and one less relevant family types of interest identified, such community members tend have more unidentified (irrelevant; i.e., individual score of "0") malware samples and families (e.g., `vendor 7`, `vendor 21`, `vendor 59`, etc.). On the other hand, certain contributors (although smaller in numbers), and despite having a small volume-based contribution, tend to have higher relevance score, thanks to having the very relevant family identified in their shared indicator labels (e.g., `vendor 10` , `vendor 16`, `vendor 27`, etc.).

Correctness. An assessment of the correctness of the AV indicators is depicted in Figure 4. As can be seen, vendor 4, vendor 27, and vendor 30, outperformed other vendors in this metric with a score in the 80s up to top 90s percentile range, highlighting a quality of the labels provided by those vendors corresponding to the actually learned labels of the samples provided by them. On the other hand, the majority of the rest of vendors tend to have a gap in the score between the volume-based and the correctness-based contribution measures, where the correctness-based measure is significantly lower. There are various reasons for why some vendors tend to score low for correctness despite their large (volume-based) contribution. This could be caused by them labeling samples under unknown names, mislabeling to other family types due to similarities between families, or assigning generic labels like "`trojan`", "`virus`", "`unclassified`" among other misleading labels. Examining the correctness of AV indicators also leads to a more subtle discussion about their utility. Looking closer into the label generated by some vendors, we find out that some labels are too generic in the sense that they only describe the behavior rather than name of a known malware family type, e.g., `Trojan.Win32.ServStart` vs. `Avzhan`.

Utility. To evaluate the utility of AV indicators, we give weights for three classes of malware labels: complete labels (ω_c) are based on industrially popular name, generic labels (ω_g) are based on placeholders commonly used for labeling the family such as "`generic`", "`worm`", "`trojan`", "`start`" and "`run`", and incomplete labels, (ω_i), including "`suspicious`", "`malware`", and "`unclassified`", which do not hold any meaning of a class. Similar to the strategy with relevance, we assign weights of $\omega_c = 5$, $\omega_g = 2$, and $\omega_i = 1$.

We plot the results of evaluating the utility of indicators in Figure 5. We notice that vendors such as `vendor 51`, `vendor 53`, and `vendor 59` are rated as high utility indicator providers that surpass their volume-based scores. Nevertheless, these vendors' high utility indicator is offset in Figure 6 that includes two more metrics: correctness and relevance. These additional metrics show that these vendors achieve insignificant correctness and relevance. Figure 6 captures the importance of displaying the three scores in a single plot to allow direct comparison of the various quality-based metrics. We notice that there is a clear correlation between the correctness and relevance scores. This is because we only

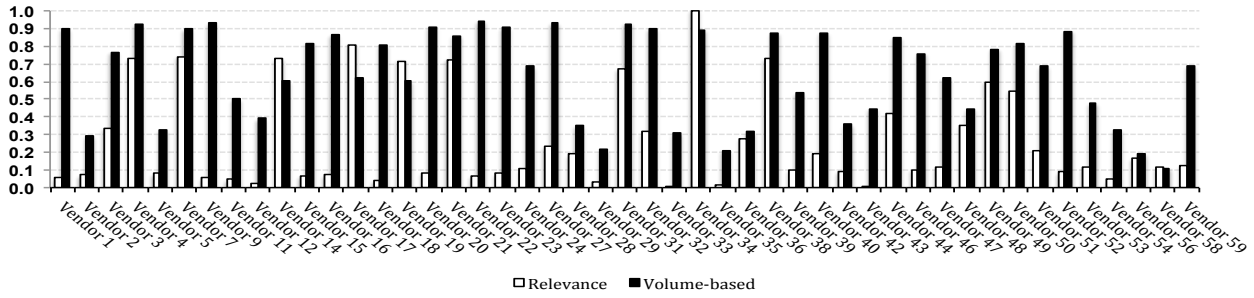

Figure 3: Comparison between relevance-based and volume-based scoring

compute relevance scores for correctly submitted samples, where incorrect labels are zeroed in the relevance score.

Aggregated QoI score. As described earlier, we aggregate a single QoI score for each vendor based on the weighted sum of the various QoI metrics. In Figure 7, we present a comparative bar-plot between QoI-based and volume-based scoring for assessment of contribution by the AV vendors. As can be seen, many vendors, such as `vendor 14`, `vendor 57`, and `vendor 58` which received high QoI scores are rated with lower scores in their volume-based rating. In particular, they received from 15-75 percent lower rating in their volume-based scores. On the other hand, vendors like `vendor 11` and `vendor 18`, and `vendor 20`, which tended to provide a very high volume-based indicators, have very small (close to zero) QoI score, highlighting their potential as free-riding candidates.

6. RELATED WORK

The free-riding problem in threat intelligence sharing is not addressed before, nor measures of quality broadly defined, or closely identified for assessing contribution by community members. However, the problem of free-riding in general is not new and has been a topic of discussion in the peer-to-peer (P2P) systems community.

Going early back in the literature, Adar and Huberman [23] were first to spot the problem in P2P systems as they noticed the existence of a large fraction of users who do not share useful content in the file sharing network Gnutella. A few years later, Feldman et al. [24] characterized the problem of free-riding in peer-to-peer systems and proposed potential directions for research. In response, a stream of papers were published on the topic, most notably, the work of Locher et al. [16], who developed a free-riding client as a proof-of-concept and demonstrated how entire files can be downloaded in the BitTorrent network without providing any content to the peers. While Locher et al. have concentrated on analyzing the feasibility of free-riding attacks, other papers [25, 26] are more focused on analyzing their root-cause and impact on the overall system utility.

In light of operationalizing the functions of threat intelligence, various information sharing standards were proposed including those developed by IETF (https://www.ietf.org/), MITRE (https://www.mitre.org/), and NIST (https://www.nist.gov/). Industry leaders have picked up on these standards and developed application program interfaces (APIs) to facilitate delivery and retrieval of raw, processed, and structure and intelligence data, such as ThreatExchange [27] by Facebook and IntelGraph by Verisign. However, sharing standards have shown to exhibit privacy violations including leaking PII fields, as demonstrated by [1], potentially encouraging the act of free-riding.

Of relevance to the notion of quality of indicators in threat intelligent systems is malware attribution. Malware attribution have been widely employed in the literature for training algorithms and techniques of malware classification and labeling [28], and understanding the utility of attributes as detector patterns of malware samples has been an important subject matter. Bailey et al. [29] were one of the early folks to characterize malware in terms of system state changes (e.g. registry changes, files created) and investigated the problem of behavior-based clustering as a method for classifying and analyzing Internet malware.

More focused on the labeling problem, Canto et al. [30] analyzed the quality of labeling of malware samples for a couple of vendors and pointed out their labeling inconsistencies. In the same vein, Perdisci [31] analyzed the shortcomings of malware labeling of various AV vendors by constructing a graph from the labels and measuring the distance between them. On the other hand, Mohaisen and Alrawi [19, 32] quantified the inconsistencies in labeling against a reference dataset collected from thousands of samples of various types which were manually vetted by analysts. In their study, the authors evaluated the detection rate, correctness, and consistency of labeling of AV scanners.

7. CONCLUSION

In this paper, we have the first look at the notion of the quality of indicators (QoI) for understanding the contribution of community members in information sharing paradigms. Unlike other peer-to-peer systems in which the volume of contribution (bandwidth, size of files, etc.) is a good indicator of contribution, we argue that the special nature of security applications calls for more elaborate notion of contribution. As such, we define multiple metrics for assessing contribution, including correctness, utility, and relevance of indicators. As compared to volume-based measures for contribution, and thus free-riding, our metrics are more robust, contextual, and reasonably quantify the actual contribution of individuals. By verifying our metrics on a real-world data of antivirus scans we unveil that contribution measured by volume is not always consistent with those quality measures, and that QoI as notion is capable of capturing forms of contribution beyond free-riding.

8. REFERENCES

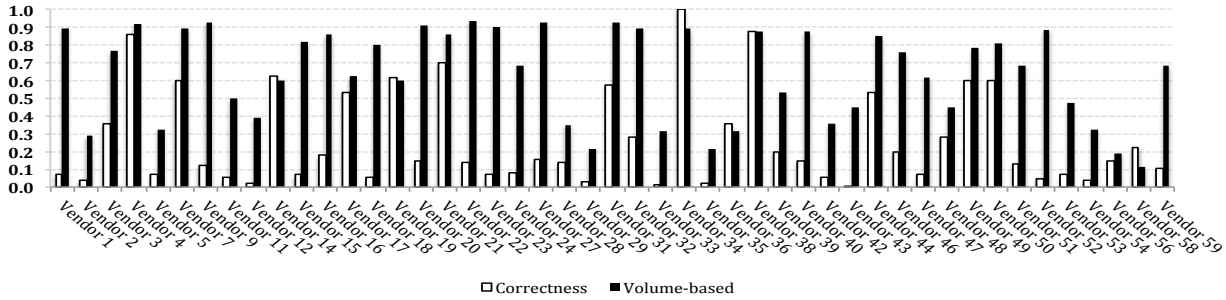

Figure 4: Comparison between correctness-based and volume-based scoring of AV indicators

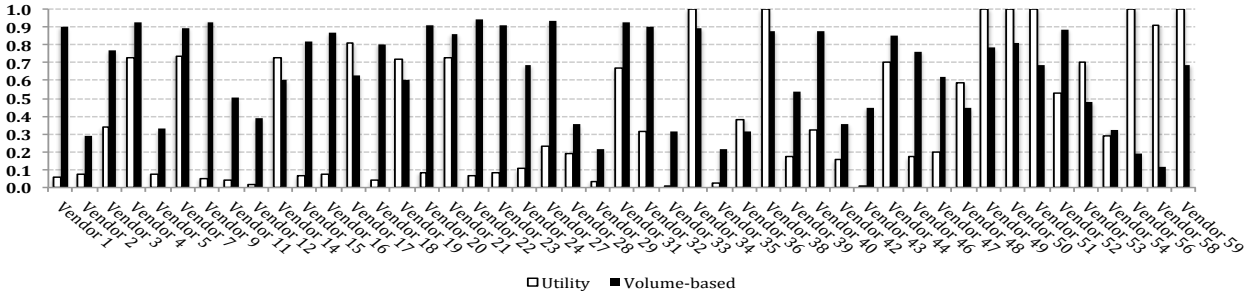

Figure 5: Comparison between utility-based and volume-based scoring

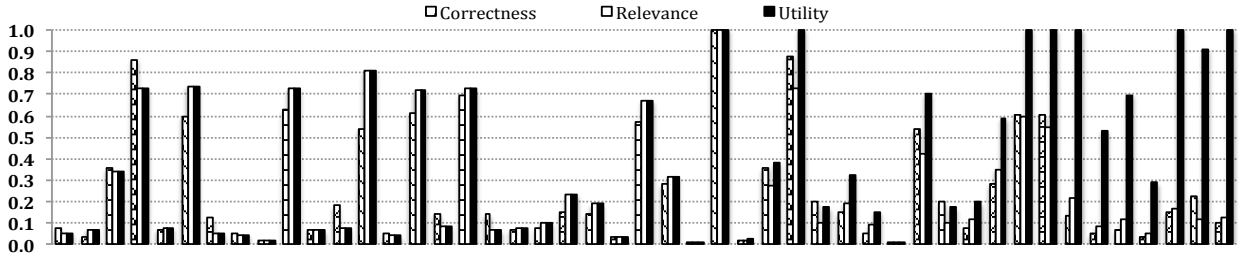

Figure 6: Comparison between various quality-based metrics for AV indicator assessment

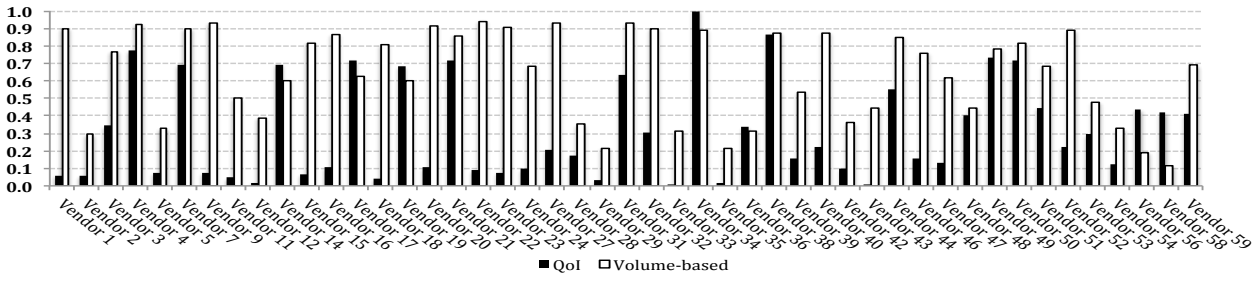

Figure 7: Comparison between QoI-based and volume-based scoring

[1] A. Mohaisen, O. Alibrahim, C. Kamhoua, K. Kwiat, and L. Njila, "Rethinking information sharing for actionable threat intelligence," tech. rep., 2016.

[2] D. K. Tosh, S. Sengupta, C. A. Kamhoua, K. A. Kwiat, and A. P. Martin, "An evolutionary game-theoretic framework for cyber-threat information sharing," in *2015 IEEE International Conference on Communications, ICC 2015, London, United Kingdom, June 8-12, 2015*, pp. 7341–7346, 2015.

[3] R. A. Martin, "Making security measurable and manageable," in *MILCOM 2008 - 2008 IEEE Military Communications Conference*, pp. 1–9, Nov 2008.

[4] S. Barnum, "Standardizing cyber threat intelligence information with the structured threat information expression (stix^{TM})," *MITRE Corporation*, vol. 11, 2012.

[5] P. Kampanakis, "Security automation and threat information-sharing options," *IEEE Security & Privacy*, vol. 12, no. 5, pp. 42–51, 2014.

[6] M. Feldman and J. Chuang, "Overcoming free-riding behavior in peer-to-peer systems," *ACM sigecom exchanges*, vol. 5, no. 4, pp. 41–50, 2005.

[7] H. Tanaka, K. Matsuura, and O. Sudoh, "Vulnerability and information security investment: An empirical analysis of e-local government in japan," *Journal of Accounting and Public Policy*, vol. 24, no. 1, pp. 37–59, 2005.

[8] D. K. Tosh, S. Sengupta, S. Mukhopadhyay, C. A. Kamhoua, and K. A. Kwiat, "Game theoretic modeling to enforce security information sharing among firms," in *IEEE 2nd International Conference on Cyber Security and Cloud Computing, CSCloud 2015, New York, NY, USA, November 3-5, 2015*, pp. 7–12, 2015.

[9] H. Mekky, A. Mohaisen, and Z.-L. Zhang, "Separation of benign and malicious network events for accurate malware family classification," in *Proc. of IEEE CNS*, 2015.

[10] A. Wang, A. Mohaisen, W. Chang, and S. Chen, "Revealing ddos attack dynamics behind the scenes," in *Proc. of DIMVA*, 2015.

[11] A. Mohaisen, "Towards automatic and lightweight detection and classification of malicious web contents," in *Proc. of IEEE HotWeb*, 2015.

[12] A. Wang, A. Mohaisen, W. Chang, and S. Chen, "Delving into internet ddos attacks by botnets: Characterization and analysis," in *Proc. of IEEE DSN*, 2015.

[13] M. Thomas and A. Mohaisen, "Kindred domains: detecting and clustering botnet domains using DNS traffic," in *Proc. of ACM WWW*, 2014.

[14] S. Yadav, A. K. K. Reddy, A. Reddy, and S. Ranjan, "Detecting algorithmically generated malicious domain names," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pp. 48–61, ACM, 2010.

[15] M. Friedewald, E. Vildjiounaite, Y. Punie, and D. Wright, "Privacy, identity and security in ambient intelligence: A scenario analysis," *Telematics and Informatics*, vol. 24, no. 1, pp. 15–29, 2007.

[16] T. Locher, P. Moor, S. Schmid, and R. Wattenhofer, "Free riding in bittorrent is cheap," in *Proc. Workshop on Hot Topics in Networks (HotNets)*, 2006.

[17] J. E. Jackson, *A user's guide to principal components*, vol. 587. John Wiley & Sons, 2005.

[18] P. C. Mahalanobis, "On the generalized distance in statistics," *Proceedings of the National Institute of Sciences (Calcutta)*, vol. 2, pp. 49–55, 1936.

[19] A. Mohaisen and O. Alrawi, *Detection of Intrusions and Malware, and Vulnerability Assessment: 11th International Conference, DIMVA 2014, Egham, UK, July 10-11, 2014. Proceedings*, ch. AV-Meter: An Evaluation of Antivirus Scans and Labels, pp. 112–131. Cham: Springer International Publishing, 2014.

[20] A. Mohaisen and O. Alrawi, "Unveiling zeus: automated classification of malware samples," in *Proc. of ACM WWW*, 2013.

[21] A. Mohaisen and O. Alrawi, "AMAL: high-fidelity, behavior-based automated malware analysis and classification," in *Proc. of WISA*, 2014.

[22] A. J. Izenman, "Linear discriminant analysis," in *Modern Multivariate Statistical Techniques*, pp. 237–280, Springer, 2013.

[23] E. Adar and B. A. Huberman, "Free riding on gnutella," *First monday*, vol. 5, no. 10, 2000.

[24] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica, "Free-riding and whitewashing in peer-to-peer systems," in *Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems*, pp. 228–236, ACM, 2004.

[25] H. Varian, "System reliability and free riding," in *Economics of information security*, pp. 1–15, Springer, 2004.

[26] D. Hughes, G. Coulson, and J. Walkerdine, "Free riding on gnutella revisited: the bell tolls?," *IEEE distributed systems online*, vol. 6, no. 6, 2005.

[27] F. Inc., "Threatexchange."

[28] C. Rossow, C. J. Dietrich, C. Grier, C. Kreibich, V. Paxson, N. Pohlmann, H. Bos, and M. Van Steen, "Prudent practices for designing malware experiments: Status quo and outlook," in *2012 IEEE Symposium on Security and Privacy*, pp. 65–79, IEEE, 2012.

[29] M. Bailey, J. Oberheide, J. Andersen, Z. M. Mao, F. Jahanian, and J. Nazario, "Automated classification and analysis of internet malware," in *International Workshop on Recent Advances in Intrusion Detection*, pp. 178–197, Springer, 2007.

[30] J. Canto, M. Dacier, E. Kirda, and C. Leita, "Large scale malware collection: lessons learned," in *IEEE SRDS Workshop on Sharing Field Data and Experiment Measurements on Resilience of Distributed Computing Systems*, Citeseer, 2008.

[31] R. Perdisci *et al.*, "Vamo: towards a fully automated malware clustering validity analysis," in *Proceedings of the 28th Annual Computer Security Applications Conference*, pp. 329–338, ACM, 2012.

[32] A. Mohaisen, O. Alrawi, and M. Mohaisen, "Amal: High-fidelity, behavior-based automated malware analysis and classification," *Computers & Security*, pp. –, 2015.